

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Manifest-Based Trusted Agent Management in a
Trusted Operating System Environment**

Inventor(s):

**Paul England
Marcus Peinado
Dan Simon
Josh Benaloh**

ATTORNEY'S DOCKET NO. MS1-953US

1

2 **TECHNICAL FIELD**

3 This invention relates to trusted environments generally, and more
4 particularly to manifest-based trusted agent management in a trusted operating
5 system environment.

6

7 **BACKGROUND**

8 Having people be able to trust computers has become an increasingly
9 important goal. This trust generally focuses on the ability to trust the computer to
10 use the information it stores or receives correctly. Exactly what this trust entails
11 can vary based on the circumstances. For example, multimedia content providers
12 would like to be able to trust computers to not improperly copy their content. By
13 way of another example, users would like to be able to trust their computers to
14 forward confidential financial information (e.g., bank account numbers) only to
15 appropriate destinations (e.g., allow the information to be passed to their bank, but
16 nowhere else). Unfortunately, given the generally open nature of most computers,
17 a wide range of applications can be run on most current computers without the
18 user's knowledge, and these applications can compromise this trust (e.g., forward
19 the user's financial information to some other destination for malicious use).

20 To address these trust issues, different mechanisms have been proposed
21 (and new mechanisms are being developed) that allow a computer or portions
22 thereof to be trusted. Generally, these mechanisms entail some sort of
23 authentication procedure where the computer can authenticate or certify that at
24 least a portion of it (e.g., certain areas of memory, certain applications, etc.) are at

1 least as trustworthy as they present themselves to be (e.g., that the computer or
2 application actually is what it claims to be). In other words, these mechanisms
3 prevent a malicious application from impersonating another application (or
4 allowing a computer to impersonate another computer). Once such a mechanism
5 can be established, the user or others (e.g., content providers) can make a
6 judgment as to whether or not to accept a particular platform and application as
7 trustworthy (e.g., a multimedia content provider may accept a particular
8 application as being trustworthy, once the computer can certify to the content
9 provider's satisfaction that the particular application is the application it claims to
10 be).

11 Oftentimes, components and modules of an application are allowed to be
12 changed (e.g., in response to user preferences) and/or upgraded fairly frequently.
13 For example, applications frequently include various dynamic link libraries
14 (DLL's), plug-ins, etc. and allow for different software configurations, each of
15 which can alter the binaries which execute as the application. Currently, it is
16 difficult (if possible at all) in many systems to allow for such changes and
17 differing configurations of applications, while at the same time maintaining the
18 trustworthiness of the computer. Thus, it would be beneficial to have a security
19 model that allows for these differences and changes, while at the same time
20 maintaining the trustworthiness of the computer. The manifest-based trusted agent
21 management in a trusted operating system environment described herein provides
22 such a security model.

1 **SUMMARY**

2 Manifest-based trusted agent management in a trusted operating system
3 environment is described herein.

4 According to one aspect, a request to execute a process is received and a
5 virtual memory space for the process is set up. A manifest corresponding to the
6 process is accessed, and which of a plurality of binaries can be executed in the
7 virtual memory space is limited based on indicators, of the binaries, that are
8 included in the manifest.

9 According to another aspect, a manifest includes a first portion including
10 data representing a unique identifier of the trusted application, a second portion
11 including data indicating whether a particular one or more binaries can be loaded
12 into the process space for the trusted application, and a third portion derived from
13 the data in both the first portion and the second portion by generating a digital
14 signature over the first and second portions. The manifest can also include a
15 portion that includes data representing a list of one or more export statements that
16 allow a secret associated with the trusted application to be exported to another
17 trusted application, a portion that includes data representing a set of properties
18 corresponding to the data structure, and a portion that includes data representing a
19 list of entry points into the executing trusted application.

20 **BRIEF DESCRIPTION OF THE DRAWINGS**

21 Fig. 1 illustrates an exemplary trusted operating system environment.

22 Fig. 2 illustrates one exemplary architecture that can be implemented on a
23 client computing device.

1 Fig. 3 illustrates another exemplary architecture that can be used with the
2 invention.

3 Fig. 4 illustrates an exemplary relationship between a gatekeeper storage
4 key and trusted application secrets.

5 Fig. 5 illustrates an exemplary process for securely storing secrets using a
6 gatekeeper storage key.

7 Fig. 6 illustrates an exemplary upgrade from one trusted core to another
8 trusted core on the same client computing device.

9 Fig. 7 illustrates an exemplary process for upgrading a trusted core.

10 Fig. 8 illustrates another exemplary process for upgrading a trusted core.

11 Fig. 9 illustrates an exemplary secret storage architecture employing hive
12 keys.

13 Fig. 10 illustrates an exemplary process for securely storing secrets using
14 hive keys.

15 Fig. 11 illustrates an exemplary process for migrating secrets from a source
16 computing device to a destination computing device.

17 Fig. 12 illustrates an exemplary manifest corresponding to a trusted
18 application.

19 Fig. 13 illustrates an exemplary process for controlling execution of
20 processes in an address space based on a manifest.

21 Fig. 14 illustrates an exemplary process for upgrading to a new version of a
22 trusted application.

23 Fig. 15 illustrates a general exemplary computer environment, which can be
24 used to implement various devices and processes described herein.

25

1

2 **DETAILED DESCRIPTION**

3 As used herein, code being "trusted" refers to code that is immutable in
4 nature and immutable in identity. Code that is trusted is immune to being
5 tampered with by other parts (e.g. code) of the computer and it can be reliably and
6 unambiguously identified. In other words, any other entity or component asking
7 "who is this code" can be told "this is code xyz", and can be assured both that the
8 code is indeed code xyz (rather than some imposter) and that code xyz is
9 unadulterated. Trust does not deal with any quality or usefulness aspects of the
10 code – only immutability of nature and immutability of identity.

11 Additionally, the execution environment of the trusted code effects the
12 overall security. The execution environment includes the machine or machine
13 class on which the code is executing.

14

15 **General Operating Environment**

16 Fig. 1 illustrates an exemplary trusted operating system environment 100.
17 In environment 100, multiple client computing devices 102 are coupled to multiple
18 server computing devices 104 via a network 106. Network 106 is intended to
19 represent any of a wide variety of conventional network topologies and types
20 (including wired and/or wireless networks), employing any of a wide variety of
21 conventional network protocols (including public and/or proprietary protocols).
22 Network 106 may include, for example, the Internet as well as possibly at least
23 portions of one or more local area networks (LANs).

1 Computing devices 102 and 104 can each be any of a wide variety of
2 conventional computing devices, including desktop PCs, workstations, mainframe
3 computers, Internet appliances, gaming consoles, handheld PCs, cellular
4 telephones, personal digital assistants (PDAs), etc. One or more of devices 102
5 and 104 can be the same types of devices, or alternatively different types of
6 devices.

7 Each of client computing devices 102 includes a secure operating system
8 (OS) 108. Secure operating system 108 is designed to provide a level of trust to
9 users of client devices 102 as well as server devices 104 that are in communication
10 with client devices 102 via a network 106. Secure operating system 108 can be
11 designed in different ways to provide such trust, as discussed in more detail below.
12 By providing this trust, the user of device 102 and/or the server devices 104 can be
13 assured that secure operating system 108 will use data appropriately and take
14 various measures to protect that data.

15 Each of client computing devices 102 may also execute one or more trusted
16 applications (also referred to as trusted agents or processes) 110. Each trusted
17 application is software (or alternatively firmware) that is made up of multiple
18 instructions to be executed by a processor(s) of device 102. Oftentimes a trusted
19 application is made up of multiple individual files (also referred to as binaries) that
20 together include the instructions that comprise the trusted application.

21 One example of the usage of environment 100 is to maintain rights to
22 digital content, often referred to as "digital rights management". A client device
23 102 may obtain digital content (e.g., a movie, song, electronic book, etc.) from a
24 server device 104. Secure operating system 108 on client device 102 assures

1 server device 104 that operating system 108 will not use the digital content
2 inappropriately (e.g., will not communicate copies of the digital content to other
3 devices) and will take steps to protect the digital content (e.g., will not allow
4 unauthorized applications to access decrypted content).

5 Another example of the usage of environment 100 is for electronic
6 commerce (also referred to as e-commerce). A client device 102 may
7 communicate with a server device 104 and exchange confidential financial
8 information (e.g., to purchase or sell a product or service, to perform banking
9 operations such as withdrawal or transfer of funds, etc.). Secure operating system
10 108 on the client device 102 assures server device 104, as well as the user of client
11 device 102, that it will not use the financial information inappropriately (e.g., will
12 not steal account numbers or funds) and will take steps to protect the financial
13 information (e.g., will not allow unauthorized applications to access decrypted
14 content).

15 Secure operating system 108 may be employed to maintain various secrets
16 by different trusted applications 110 executing on client devices 102. For
17 example, confidential information may be encrypted by a trusted application 110
18 and a key used for this encryption securely stored by secure operating system 108.
19 By way of another example, the confidential information itself may be passed to
20 secure operating system 108 for secure storage.

21 There are two primary functions that secure operating system 108 provides:
22 (1) the ability to securely store secrets for trusted applications 110; and (2) the
23 ability to allow trusted applications 110 to authenticate themselves. The secure
24 storage of secrets allows trusted applications 110 to save secrets to secure
25

1 operating system 108 and subsequently retrieve those secrets so long as neither the
2 trusted application 110 nor operating system 108 has been altered. If either the
3 trusted application 110 or the operating system 108 has been altered (e.g., by a
4 malicious user or application in an attempt to subvert the security of operating
5 system 108), then the secrets are not retrievable by the altered application and/or
6 operating system. A secret refers to any type of data that the trusted application
7 does not want to make publicly available, such as an encryption key, a user
8 password, a password to access a remote computing device, digital content (e.g., a
9 movie, a song, an electronic book, etc.) or a key(s) used to encrypt the digital
10 content, financial data (e.g., account numbers, personal identification numbers
11 (PINs), account balances, etc.), and so forth.

12 The ability for a trusted application 110 to authenticate itself allows the
13 trusted application to authenticate itself to a third party (e.g., a server device 104).
14 This allows, for example, a server device 104 to be assured that it is
15 communicating digital content to a trusted content player executing on a trusted
16 operating system, or for the server device 104 to be assured that it is
17 communicating with a trusted e-commerce application on the client device rather
18 than with a virus (or some other malicious or untrusted application).

19 Various concerns exist for the upgrading, migrating, and backing up of
20 various components of the client devices 102. As discussed in more detail below,
21 the security model discussed herein provides for authentication and secret storage

1 in a trusted operating system environment, while at the same time allowing one or
2 more of:

- 3 • secure operating system upgrades
- 4 • migration of secrets to other computing devices
- 5 • backup of secrets
- 6 • trusted application upgrades

7 Reference is made herein to encrypting data using a key. Generally,
8 encryption refers to a process in which the data to be encrypted (often referred to
9 as plaintext) is input to an encryption algorithm that operates, using a key
10 (commonly referred to as the encryption key), on the plaintext to generate
11 ciphertext. Encryption algorithms are designed so that it is extremely difficult to
12 re-generate the plaintext without knowing a decryption key (which may be the
13 same as the encryption key, or alternatively a different key). A variety of
14 conventional encryption algorithms can be used, such as DES (Data Encryption
15 Standard), RSA (Rivest, Shamir, Adelman), RC4 (Rivest Cipher 4), RC5 (Rivest
16 Cipher 5), etc.

17 One type of encryption uses a public-private key pair. The public-private
18 key pair includes two keys (one private key and one public key) that are selected
19 so that it is relatively straight-forward to decrypt the ciphertext if both keys are
20 known, but extremely difficult to decrypt the ciphertext if only one (or neither) of
21 the keys is known. Additionally, the encryption algorithm is designed and the
22 keys selected such that it is extremely difficult to determine one of the keys based
23 on the ciphertext alone and/or only one key.

The owner of a public-private key pair typically makes its public key publicly available, but keeps its private key secret. Any party or component desiring to encrypt data for the owner can encrypt the data using the owner's public key, thus allowing only the owner (who possesses the corresponding private key) to readily decrypt the data. The key pair can also be used for the owner to digitally sign data. In order to add a digital signature to data, the owner encrypts the data using the owner's private key and makes the resultant ciphertext available with the digitally signed data. A recipient of the digitally signed data can decrypt the ciphertext using the owner's public key and compare the decrypted data to the data sent by the owner to verify that the owner did in fact generate that data (and that it has not been altered by the owner since being generated).

The discussions herein assume a basic understanding of cryptography. For a basic introduction of cryptography, the reader is directed to a text written by Bruce Schneier and entitled "Applied Cryptography: Protocols, Algorithms, and Source Code in C," published by John Wiley & Sons with copyright 1994 (or second edition with copyright 1996).

Exemplary Computing Device Architectures

Secure operating system 108 of Fig. 1 includes at least a portion that is trusted code, referred to as the "trusted core". The trusted core may be a full operating system, a microkernel, a Hypervisor, or some smaller component that provides specific security services.

Fig. 2 illustrates one exemplary architecture that can be implemented on a client computing device 102. In Fig. 2, the trusted core is implemented by taking

1 advantage of different privilege levels of the processor(s) of the client computing
2 device 102 (e.g., referred to as "rings" in an x86 architecture processor). In the
3 illustrated example, these privilege levels are referred to as rings, although
4 alternate implementations using different processor architectures may use different
5 nomenclature. The multiple rings provide a set of prioritized levels that software
6 can execute at, often including 4 levels (Rings 0, 1, 2, and 3). Ring 0 is typically
7 referred to as the most privileged ring. Software processes executing in Ring 0
8 can typically access more features (e.g., instructions) than processes executing in
9 less privileged rings.

10 Furthermore, a processor executing in a particular ring cannot alter code or
11 data in a higher priority ring. In the illustrated example, a trusted core 120
12 executes in Ring 0, while an operating system 122 executes in Ring 1 and trusted
13 applications 124 execute in Ring 3. Thus, trusted core 120 operates at a more
14 privileged level and can control the execution of operating system 122 from this
15 level. Additionally, the code and/or data of trusted core 120 (executing in Ring 0)
16 cannot be altered directly by operating system 122 (executing in Ring 1) or trusted
17 applications 124 (executing in Ring 3). Rather, any such alterations would have to
18 be made by the operating system 122 or a trusted application 124 requesting
19 trusted core 120 to make the alteration (e.g., by sending a message to trusted core
20 120, invoking a function of trusted core 120, etc.).

21 Trusted core 120 also maintains a secret store 126 where secrets passed to
22 and encrypted by trusted core 120 (e.g., originating with trusted applications 124,
23 OS 122, or trusted core 120) are securely stored. The storage of secrets is
24 discussed in more detail below.

1 A cryptographic measure of trusted core 120 is also generated when it is
2 loaded into the memory of computing device 102 and stored in a digest register of
3 the hardware. In one implementation, the digest register is designed to be written
4 to only once after each time the computing device is reset, thereby preventing a
5 malicious user or application from overwriting the digest of the trusted core. This
6 cryptographic measure can be generated by different components, such as a
7 security processor of computing device 102, a trusted BIOS, etc. The
8 cryptographic measure provides a small (relative to the size of the trusted core)
9 measure of the trusted core that can be used to verify the trusted core that is
10 loaded. Given the nature of the cryptographic measure, it is most likely that any
11 changes made to a trusted core (e.g., to circumvent its trustworthiness) will be
12 reflected in the cryptographic measure, so that the altered core and the original
13 core will produce different cryptographic measures. This cryptographic measure
14 is used as a basis for securely storing data, as discussed in more detail below.

15 A variety of different cryptographic measures can be used. One such
16 cryptographic measure is a digest – for ease of explanation the cryptographic
17 measure will be discussed primarily herein as a digest, although other measures
18 could alternatively be used. The digest is calculated using a one-way hashing
19 operation, such as SHA-1 (Secure Hash Algorithm 1), MD4 (Message Digest 4),
20 MD5 (Message Digest 5), etc. The cryptographic digest has the property that it is
21 extremely difficult to find a second pre-image (in this case, a second trusted core)
22 that when digested produces the same hash value. Hence the digest register
23 contains a value that can be considered to uniquely represent the trusted core in
24 use.

1 An alternative cryptographic measure to a digest, is the public key of a
2 properly formed certificate on the digest. Using this technique, a publisher can
3 generate a sequence of trusted-cores that are treated as identical or equivalent by
4 the platform (e.g., based on the public key of the publisher). The platform refers
5 to the basic hardware of the computing device (e.g., processor and chipset) as well
6 as the firmware associated with this hardware (e.g., microcode in the processor
7 and/or chipset).

8 Alternatively, the operating system may be separated into a memory
9 manager component that operates as trusted core 120 with the remainder of the
10 operating system operating as OS 122. The trusted core 120 then controls all page
11 maps and is thus able to shield trusted agents executing in Ring 3 from other
12 components (including OS 122). In this alternative, additional control is also
13 added to protect the trusted core 120 from other busmasters that do not obey ring
14 privileges.

15 Fig. 3 illustrates another exemplary architecture that can be used with the
16 invention. In Fig. 3, the trusted core is implemented by establishing two separate
17 "spaces" within a client computing device 102 of Fig. 1: a trusted space 140 (also
18 referred to as a protected parallel area, or curtained memory) and a normal
19 (untrusted) space 142. These spaces can be, for example, one or more address
20 ranges within computing device 102. Both trusted space 140 and normal space
21 142 include a user space and a kernel space, with the trusted core 144 and secret
22 store 146 being implemented in the kernel space of trusted space 140. A
23 cryptographic measure, such as a digest, of trusted core 144 is also generated and
24 used analogous to the cryptographic measure of trusted core 120 discussed above.

1 A variety of trusted applets, trusted applications, and/or trusted agents 148
2 can execute within the user space of trusted space 140, under the control of trusted
3 core 144. However, any application 150, operating system 152, or device driver
4 executing in normal space 142 is prevented, by trusted core 144, from
5 accessing trusted space 140. Thus, no alterations can be made to trusted
6 applications or data in trusted space 140 unless approved by trusted core 144.

7 Additional information regarding these computing device architectures can
8 be found in the following four U.S. Patent Applications, each of which is hereby
9 incorporated by reference: U.S. Patent Application No. 09/227,611, entitled
10 "Loading and Identifying a Digital Rights Management Operating System", which
11 was filed January 8, 1999, in the names of Paul England et al.; U.S. Patent
12 Application No. 09/227,561, entitled "Digital Rights Management Operating
13 System", which was filed January 8, 1999, in the names of Paul England et al.;
14 U.S. Patent Application No. 09/287,393, entitled "Secure Execution of Program
15 Code", which was filed April 6, 1999, in the names of Paul England et al.; and
16 U.S. Patent Application No. 09/287,698, entitled "Hierarchical Trusted Code for
17 Content Protection in Computers", which was filed April 6, 1999, in the name of
18 Paul England.

19 For ease of explanation, the digest of a trusted core is discussed herein as a
20 single digest of the trusted core. However, in different implementations, the digest
21 may be made up of multiple parts. By way of example, the boot process may
22 involve a trusted BIOS loading a platform portion of the trusted core and
23 generating a digest of the platform portion. The platform portion in turn loads an
24 operating system portion of the trusted core and generates a digest for the

1 operating system portion. The operating system portion in turn loads a gatekeeper
2 portion of the trusted core and generates a digest for the gatekeeper portion. A
3 composite of these multiple generated digests is used as the digest of the trusted
4 core. These multiple generated digests may be stored individually in separate
5 digest registers with the composite of the digests being the concatenation of the
6 different register values. Alternatively, each new digest may be used to generate a
7 new digest value by generating a cryptographic hash of the previous digest value
8 concatenated with the new digest – the last new digest value generated (e.g., by
9 the operating system portion) is stored in a single digest register.

10

11 Primitive Operations

12 Two fundamental types of primitive operations are supported by the
13 hardware and software of a client computing device 102 of Fig. 1. These
14 fundamental types are secret storage primitives and authentication primitives. The
15 hardware of a device 102 makes these primitive operations available to the trusted
16 core executing on the device 102, and the trusted core makes variations of these
17 primitive operations available to the trusted applications executing on the device
18 102.

19 Two secret storage primitive operations are supported: Seal and Unseal.
20 The Seal primitive operation uses at least two parameters – one parameter is the
21 secret that is to be securely stored and the other parameter is an identification of
22 the module or component that is to be able to subsequently retrieve the secret. In
23 one implementation, the Seal primitive operation provided by the hardware of

1 client computing device 102 (e.g., by a cryptographic or security processor of
2 device 102) takes the following form:

3 Seal (*secret*, *digest_to_unseal*, *current_digest*)

4 where *secret* represents the secret to be securely stored, *digest_to_unseal*
5 represents a cryptographic digest of the trusted core that is authorized to
6 subsequently retrieve the secret, and *current_digest* represents a cryptographic
7 digest of the trusted core at the time the Seal operation was invoked. The
8 *current_digest* the *current_digest* is automatically added by the security processor
9 as the value in the digest register of the device 102 rather than being explicitly
10 settable as an external parameter (thereby removing the possibility that the module
11 or component invoking the Seal operation provides an inaccurate *current_digest*).

12 When the Seal primitive operation is invoked, the security processor
13 encrypts the parameters provided (e.g., *secret*, *digest_to_unseal*, and
14 *current_digest*). Alternatively, the *digest_to_unseal* (and optionally the
15 *current_digest* as well) may not be encrypted, but rather stored in non-encrypted
16 form and a correspondence maintained between the encrypted secret and the
17 *digest_to_unseal*. By not encrypting the *digest_to_unseal*, comparisons performed
18 in response to the Unseal primitive operation discussed below can be carried out
19 without decrypting the ciphertext.

20 The security processor can encrypt the data of the Seal operation in any of a
21 wide variety of conventional manners. For example, the security processor may
22 have an individual key that it keeps secret and divulges to no component or
23 module, and/or a public-private key pair. The security processor could use the
24 individual key, the public key from its public-private key pair, or a combination

1 thereof. The security processor can use any of a wide variety of conventional
2 encryption algorithms to encrypt the data. The resultant ciphertext is then stored
3 as a secret (e.g., in secret store 126 of Fig. 2 or 146 of Fig. 3).

4 The Unseal primitive operation is the converse of the Seal primitive
5 operation, and takes as a single parameter the ciphertext produced by an earlier
6 Seal operation. The security processor obtains the cryptographic digest of the
7 trusted core currently executing on the computing device and also obtains the
8 *digest_to_unseal*. If the *digest_to_unseal* exists in a non-encrypted state (e.g.,
9 associated with the ciphertext, but not encrypted as part of the ciphertext), then
10 this non-encrypted version of the *digest_to_unseal* is obtained by the security
11 processor. However, if no such non-encrypted version of the *digest_to_unseal*
12 exists, then the security processor decrypts the ciphertext to obtain the
13 *digest_to_unseal*.

14 Once the *digest_to_unseal* and the cryptographic digest of the trusted core
15 currently executing on the computing device are both obtained, the security
16 processor compares the two digests to determine if they are the same. If the two
17 digests are identical, then the trusted core currently executing on the computing
18 device is authorized to retrieve the secret, and the security processor returns the
19 secret (decrypting the secret, if it has not already been decrypted) to the
20 component or module invoking the Unseal operation. However, if the two digests
21 are not identical, then the trusted core currently executing on the computing device
22 is not authorized to retrieve the secret and the security processor does not return
23 the secret (e.g., returning a "fail" notification). Note that failures of the Unseal
24 operation will also occur if the ciphertext was generated on a different platform
25

1 (e.g., a computing device using a different platform firmware) using a different
2 encryption or integrity key, or if the ciphertext was generated by some other
3 process (although the security processor may decrypt the secret and make it
4 available to the trusted core, the trusted core would not return the secret to the
5 other process).

6 Two authentication primitive operations are also supported: Quote and
7 Unwrap (also referred to as PK_Unseal). The Quote primitive takes one
8 parameter, and causes the security processor to generate a signed statement
9 associating the supplied parameter with the digest of the currently running trusted
10 core. In one implementation, the security processor generates a certificate that
11 includes the public key of a public-private key pair of the security processor as
12 well as the digest of the currently running trusted core and the external parameter.
13 The security processor then digitally signs this certificate and returns it to the
14 component or module (and possibly ultimately to a remote third party), which can
15 use the public key in the certificate to verify the signature.

16 The Unwrap or PK_Unseal primitive operation, has ciphertext as its single
17 parameter. The party invoking the Unwrap or PK_Unseal operation initially
18 generates a structure that includes two parts – a *secret* and a *digest_to_unseal*.
19 The party then encrypts this structure using the public key of a public-private key
20 pair of the security processor on the client computing device 102. The security
21 processor responds to the Unwrap or PK_Unseal primitive operation by using its
22 private key of the public-private key pair to decrypt the ciphertext received from
23 the invoking party. Similar to the Unseal primitive operation discussed above, the
24 security processor compares the digest of the trusted core currently running on the
25

1 client computing device 102 to the *digest_to_unseal* from the decrypted
2 ciphertext. If the two digests are identical, then the trusted core currently
3 executing on the computing device is authorized to retrieve the secret, and the
4 security processor provides the secret to the trusted core. However, if the two
5 digests are not identical, then the trusted core currently executing on the
6 computing device is not authorized to retrieve the secret and the security processor
7 does not provide the secret to the trusted core (e.g., instead providing a "fail"
8 notification).

9 Both quote and unwrap can be used as part of a cryptographic protocol that
10 allows a remote party to be assured that he is communicating with a trusted
11 platform running a specific piece of trusted core software (by knowing its digest).
12

13 **Gatekeeper Storage Key and Trusted Core Updates**

14 Secret use and storage by trusted applications executing on a client
15 computing device 102 of Fig. 1 is based on a key generated by the trusted core,
16 referred to as the gatekeeper storage key (GSK). The gatekeeper storage key is
17 used to facilitate upgrading of the secure part of the operating system (the trusted
18 core) and also to reduce the frequency with which the hardware Seal primitive
19 operation is invoked. The gatekeeper storage key is generated by the trusted core
20 and then securely stored using the Seal operation with the digest of the trusted core
21 itself being the *digest_to_unseal* (this is also referred to as sealing the gatekeeper
22 storage key to the trusted core with the digest *digest_to_unseal*). Securely storing
23 the gatekeeper storage key using the Seal operation allows the trusted core to
24 retrieve the gatekeeper storage key when the trusted core is subsequently re-
25

1 booted (assuming that the trusted core has not been altered, and thus that its digest
2 has not been altered). The trusted core should not disclose the GSK to any other
3 parties, apart from under the strict rules detailed below.

4 The gatekeeper storage key is used as a root key to securely store any
5 trusted application, trusted core, or other operating system secrets. A trusted
6 application desiring to store data as a secret invokes a software implementation of
7 Seal supported by the trusted core (e.g., exposed by the trusted core via an
8 application programming interface (API)). The trusted core encrypts the received
9 trusted application secret using an encryption algorithm that uses the gatekeeper
10 storage key as its encryption key. Any of a wide variety of conventional
11 encryption algorithms can be used. The encrypted secret is then stored by the
12 trusted core (e.g., in secret store 126 of Fig. 2, secret store 146 of Fig. 3, or
13 alternatively elsewhere (typically, but not necessarily, on the client device)).

14 When the trusted application desires to subsequently retrieve the stored
15 secret, the trusted application invokes an Unseal operation supported by the trusted
16 core (e.g., exposed by the trusted core via an API) and based on the GSK as the
17 encryption key. The trusted core determines whether to allow the trusted
18 application to retrieve the secret based on information the trusted core has about
19 the trusted application that saved the secret as well as the trusted application that is
20 requesting the secret. Retrieval of secrets is discussed in more detail below with
21 reference to manifests.

22 Thus, the gatekeeper storage key allows multiple trusted application secrets
23 to be securely stored without the Seal operation of the hardware being invoked a
24 corresponding number of times. However, security of the trusted application
25

1 secrets is still maintained because a mischievous trusted core will not be able to
2 decrypt the trusted application secrets (it will not be able to recover the gatekeeper
3 storage key that was used to encrypt the trusted application secrets, and thus will
4 not be able to decrypt the encrypted trusted application secrets).

5 Fig. 4 illustrates an exemplary relationship between the gatekeeper storage
6 key and trusted application secrets. A single gatekeeper storage key 180 is a root
7 key and multiple (n) trusted application secrets 182, 184, and 186 are securely
8 stored based on key 180. Trusted application secrets 182, 184, and 186 can be
9 stored by a single trusted application or alternatively multiple trusted applications.
10 Each trusted application secret 182, 184, and 186 optionally includes a policy
11 statement 188, 190, and 192, respectively. The policy statement includes policy
12 information regarding the storage, usage, and/or migration conditions that the
13 trusted application desires to be imposed on the corresponding trusted application
14 secret.

15 Fig. 5 illustrates an exemplary process 200 for securely storing secrets
16 using a gatekeeper storage key. The process of Fig. 5 is carried out by the trusted
17 core of a client computing device, and may be performed in software.

18 The first time the trusted core is booted, a gatekeeper storage key is
19 obtained (act 202) and optionally sealed, using a cryptographic measure of the
20 trusted core, to the trusted core (act 204). The gatekeeper storage key may not be
21 sealed, depending on the manner in which the gatekeeper storage keys are
22 generated, as discussed in more detail below. Eventually, a request to store a
23 secret is received by the trusted core from a trusted application (act 206). The
24
25

1 trusted core uses the gatekeeper storage key to encrypt the trusted application
2 secret (act 208), and stores the encrypted secret.

3 The gatekeeper storage key can be generated in a variety of different
4 manners. In one implementation, the trusted core generates a gatekeeper storage
5 key by generating a random number (or pseudo-random number) and uses a seal
6 primitive to save and protect it between reboots. This generated gatekeeper
7 storage key can also be transferred to other computing devices under certain
8 circumstances, as discussed in more detail below. In another implementation,
9 platform firmware on a computing device generates a gatekeeper storage key
10 according to a particular procedure that allows any previous gatekeeper storage
11 keys to be obtained by the trusted core, but does not allow the trusted core to
12 obtain any future gatekeeper storage keys; in this case an explicit seal/unseal step
13 need not be performed.

14 With this secret storage structure based on the gatekeeper storage key, the
15 trusted core on the client computing device may be upgraded to a new trusted core
16 and these secrets maintained. Fig. 6 illustrates an exemplary upgrade from one
17 trusted core to another trusted core on the same client computing device.

18 The initial trusted core executing on the client computing device is trusted
19 core(0) 230, which is to be upgraded to trusted core(1) 232. Trusted core 230
20 includes (or corresponds to) a certificate 234, a public key 236, and a gatekeeper
21 storage key 238 (GSK_0). Public key 236 is the public key of a public-private key
22 pair of the component or device that is the source of trusted core 230 (e.g., the
23 manufacturer of trusted core 230). Certificate 234 is digitally signed by the source
24 of trusted core 230, and includes the digest 240 of trusted core 230. Similarly,

1 trusted core 232 includes (or corresponds to) a certificate 242 including a digest
2 244, and a public key 246. After trusted core 230 is upgraded to trusted core 232,
3 trusted core 232 will also include a gatekeeper storage key 248 (GSK_1), as well as
4 gatekeeper storage key 238 (GSK_0). Optionally, trusted cores 230 and 232 may
5 also include version identifiers 250 and 252, respectively.

6 Fig. 7 illustrates an exemplary process 270 for upgrading a trusted core
7 which uses the seal/unseal primitives. The process of Fig. 7 is carried out by the
8 two trusted cores. The process of Fig. 7 is discussed with reference to components
9 of Fig. 6. For ease of explanation, the acts performed by the initial trusted core
10 (trusted core(0)) are on the left-hand side of Fig. 7 and the acts performed by the
11 new trusted core (trusted core(1)) are on the right-hand side of Fig. 7.

12 Initially, a request to upgrade trusted core(0) to trusted core(1) is received
13 (act 272). The upgrade request is accompanied by the certificate belonging to the
14 proposed upgrade trusted core (trusted core (1)). Trusted core(0) verifies the
15 digest of proposed-upgraded trusted core(1) (act 274), such as by using public key
16 246 to verify certificate 242. Trusted core(0) also optionally checks whether one
17 or more other upgrade conditions are satisfied (act 276). Any of a variety of
18 upgrade conditions may be imposed. In one implementation, trusted core(0)
19 imposes the restriction that trusted cores are upgraded in strictly increasing version
20 numbers and are signed by the same certification authority as the one that certified
21 the currently running trusted core (or alternatively signed by some other key
22 known to by the currently running trusted core to be held by a trusted publisher).
23 Thus, version 0 can only be replaced by version 1, version 1 can only be replaced
24 by version 2, and so forth. In most cases, it is also desirable to allow version 0 to
25

1 be upgraded to version 2 in a single step (e.g., without having to be upgraded to
2 version 1 in between). However, it is generally not desirable to allow
3 “downgrades” to earlier versions (e.g., earlier versions may have more security
4 vulnerabilities).

5 If the check in act 276 determines that the various conditions (including the
6 verification of the digest in act 274) are not satisfied, then the upgrade process
7 fails and the trusted core refuses to seal the gatekeeper storage key to the
8 prospective-newer trusted core (act 278). Thus, even if the prospective-newer
9 trusted core were to be installed on the computing device, it would not have access
10 to any secrets stored by trusted core (0). However, if the various conditions are
11 satisfied, then the upgrade process is authorized to proceed and trusted core(0)
12 uses the Seal primitive operation to seal gatekeeper storage key 238 to the digest
13 of trusted core(1) as stated in the certificate received in act 272 (act 280). In
14 sealing the GSK 238 to the digest of trusted core(1), trusted core(0) uses the Seal
15 operation with digest 244 being the *digest_to_unseal* parameter.

16 Once the Seal operation is completed, trusted core(1) may be loaded and
17 booted. This may be an automated step (e.g., performed by trusted core(0)), or
18 alternatively a manual step performed by a user or system administrator.

19 Once trusted core(1) is loaded and booted, trusted core(1) obtains the sealed
20 gatekeeper storage key 238 (act 282). Trusted core(1) unseals gatekeeper storage
21 key 238 (act 284), which it is able to successfully do as its digest 244 matches the
22 *digest_to_unseal* parameter used to seal gatekeeper storage key 238. Trusted
23 core(1) then generates its own gatekeeper storage key 248 (act 286) and seals
24 gatekeeper storage key 248 to the trusted core(1) digest (act 288), thereby
25

1 allowing gatekeeper storage key 248 to be subsequently retrieved by trusted
2 core(1). Trusted core (1) may also optionally seal gatekeeper storage key 238 to
3 the trusted core(1) digest. For subsequent requests by trusted applications to store
4 secrets, trusted core(1) uses gatekeeper storage key 248 to securely store the
5 secrets (act 290). For subsequent requests by trusted applications to retrieve
6 secrets, trusted core(1) uses gatekeeper storage key 238 to retrieve old secrets
7 (secrets that were sealed by trusted core(0)), and uses gatekeeper storage key 248
8 to retrieve new secrets (secrets that were sealed by trusted core(1)) (act 292).

9 Returning to Fig. 5, another way in which the gatekeeper storage key may
10 be obtained (act 200) is by having the platform generate a set of one or more keys
11 to be used as gatekeeper storage keys. By way of example, the platform can
12 generate a set of gatekeeper storage keys (*SK*) for trusted cores according to the
13 following calculation:

14 $SK_n = \text{SHA-1}(\text{cat}(BK, \text{public_key}, n))$, for $n=0$ to N

15 where *BK* is a unique platform key called a binding key which is not disclosed to
16 other parties, and is only used for the generation of keys as described above,
17 *public_key* represents the public key of the party that generated the trusted core for
18 which the gatekeeper storage keys are being generated, and *N* represents the
19 version number of the trusted core. When booting a particular trusted core “n”,
20 the platform generates the family of keys from 1 to n and provide them to trusted
21 core “n.” Each time trusted core n boots, it has access to all secrets stored with
22 key n (which is used as a GSK). But additionally, it has access to all secrets stored
23 with previous versions of the trusted core, because the platform has provided the
24 trusted core with all earlier secrets.

25

1 It should be noted, however, that the core cannot get access to secrets stored
2 by future trusted cores because trusted core “n” obtains the family of keys 1 to n
3 from the platform, but does not obtain key n+1 or any other keys beyond n.
4 Additionally, secrets available to each family of trusted cores (identified by the
5 public key of the signer of the trusted cores) are inaccessible to cores generated by
6 a different software publisher that does not have access to the private key used to
7 generate the certificates. The certificates are provided along with the trusted core
8 (e.g., shipped by the publisher along with the trusted core), allowing the platform
9 to generate gatekeeper storage keys for that publisher’s trusted cores (based on the
10 publisher’s public key).

11 Fig. 8 illustrates an exemplary process 300 for upgrading a trusted core
12 which uses the family-based set of platform-generated gatekeeper storage keys.
13 The process of Fig. 8 is carried out by the trusted core and the platform. For ease
14 of explanation, the acts performed by the trusted core are on the left-hand side of
15 Fig. 8 and the acts performed by the platform are on the right-hand side of Fig. 8.

16 Initially, trusted core (n) requests a set of keys from the platform (act 302).
17 This request is typically issued when trusted core (n) is booted. In response to the
18 request, the platform generates a set of keys from 1 to n (act 304) and returns the
19 set of keys to trusted core (n) (act 306). Trusted core (n) eventually receives
20 requests to store and/or retrieve secrets, and uses the received set of keys to store
21 and retrieve such secrets. Trusted core (n) uses key (n) as the gatekeeper storage
22 key to store and retrieve any new secrets (act 308), and uses key (n-a) as the
23 gatekeeper storage key to retrieve any old secrets stored by a previous trusted core
24 (n-a) (act 310).

1 It should be noted that the process of Fig. 8 is the process performed by a
2 trusted core when it executes, regardless of whether it is a newly upgraded-to
3 trusted core or a trusted core that has been installed and running for an extended
4 period of time. Requests to upgrade to new trusted cores can still be received and
5 upgrades can still occur with the process of Fig. 8, but sealing of a gatekeeper
6 storage key to the digest of the new trusted core need not be performed.

7 Following a successful upgrade (regardless of the manner in which
8 gatekeeper storage keys are obtained by the trusted cores), trusted core (1) has a
9 storage facility (GSK_1) that allows it to store new secrets that will be inaccessible
10 to trusted core (0), and yet still has access to the secrets stored by trusted core (0)
11 by virtue of its access to GSK_0 . Furthermore, a user can still boot the older trusted
12 core (0) and have access to secrets that it has stored, and yet not have access to
13 newer secrets obtained by, or generated by trusted core (1).

14 Alternatively, rather than a single gatekeeper storage key, multiple
15 gatekeeper storage keys may be used by a computing device. These additional
16 second-level gatekeeper storage key(s) may be used during normal operation of
17 the device, or alternatively only during the upgrade process. Using multiple
18 gatekeeper storage keys allows trusted applications to prevent their secrets from
19 being available to an upgraded trusted core. Some trusted applications may allow
20 their secrets to be available to an upgraded trusted core, whereas other trusted
21 applications may prevent their secrets from being available to the upgraded trusted
22 core. Additionally, a particular trusted application may allow some of its secrets
23 to be available to the upgraded trusted core, but not other secrets. In one
24 implementation, when a trusted application stores a secret it indicates to the
25

1 trusted core whether the secret should be accessible to an upgraded trusted core,
2 and this indication is saved as part of the policy corresponding to the secret (e.g.,
3 policy 188, 190, or 192 of Fig. 4). The family of second-level gatekeeper storage
4 keys can be generated randomly and held encrypted by the root (sealed)
5 gatekeeper storage key. During the trusted core upgrade process, only those
6 trusted application secrets that are to be accessible to an upgraded trusted core are
7 encrypted so as to be retrievable by the upgraded trusted core. For example, the
8 trusted core being upgraded can generate a temporary gatekeeper storage key and
9 encrypt a subset of the trusted application secrets (all of the secrets that are to be
10 retrievable by the upgraded trusted core) using the temporary gatekeeper storage
11 key. The temporary gatekeeper storage key is then sealed to the digest of the new
12 trusted core, but the other gatekeeper storage key used by the trusted core is not
13 sealed to the digest of the new trusted core. Thus, when the new trusted core is
14 loaded and booted, the new trusted core will be able to retrieve the temporary
15 gatekeeper storage key and thus retrieve all of the trusted application secrets that
16 were saved using the temporary gatekeeper storage key, but not trusted application
17 secrets that were saved using the other gatekeeper storage key.

18 Thus, the trusted core upgrade process allows the new upgraded trusted
19 core to access secrets that were securely stored by the previous trusted core(s), as
20 the new upgraded trusted core has access to the gatekeeper storage key used by the
21 previous trusted core(s). However, any other core (e.g., a mischievous core)
22 would not have the same digest as the new upgraded trusted core, or would not
23 have a valid certificate (digitally signed with the private key of the publisher of the
24 new upgraded trusted core) with the public key of the publisher of the new
25

1 upgraded trusted core, and thus would not have access to the secrets. Furthermore,
2 if a previous trusted core were to be loaded and executed after secrets were stored
3 by the new upgraded trusted core, the previous trusted core would not have access
4 to the secrets stored by the new upgraded trusted core because the previous trusted
5 core is not able to retrieve the gatekeeper storage key of the new upgraded trusted
6 core. Additionally, the trusted core upgrade process allows the new upgraded
7 trusted core to be authenticated to third parties. The security processor uses the
8 digest of the new upgraded trusted core in performing any Quote or
9 Unwrap/PK_Unseal primitive operations.

10

11 **Hive Keys and Secret Migration**

12 Secret use and storage by trusted applications executing on a client
13 computing device 102 of Fig. 1 can be further based on multiple additional keys
14 referred to as "hive" keys. The hive keys are used to facilitate migrating of trusted
15 application secrets from one computing device to another computing device. In
16 one implementation, up to three different types or classes of secrets can be
17 securely stored: non-migrateable secrets, user-migrateable secrets, and third
18 party-migrateable secrets. One or more hive keys may be used in a computing
19 device 102 for each type of secret. Trusted application secrets are securely stored
20 by encrypting the secrets using one of these hive keys. Which type of secret is
21 being stored (and thus which hive key to use) is identified by the trusted
22 application when storing the secret (e.g., is a parameter of the seal operation that
23 the trusted core makes available to the trusted applications). Whether a particular
24
25

1 trusted application secret can be migrated to another computing device is
2 dependent on which type of secret it is.

3 Fig. 9 illustrates an exemplary secret storage architecture employing hive
4 keys. A root gatekeeper storage key 320 and three types of hive keys are included:
5 a non-migrateable key 322, one or more user-migrateable keys 324, and one or
6 more third party-migrateable keys 326. Non-migrateable trusted application
7 secrets 328 are encrypted by the trusted core using non-migrateable key 322, user-
8 migrateable trusted application secrets 330 are encrypted by the trusted core using
9 user-migrateable key 324, and third party-migrateable secrets 332 are encrypted
10 by the trusted core using third party-migrateable key 326.

11 Each of the hive keys 322, 324, and 326, in turn, is encrypted by the trusted
12 core using gatekeeper storage key 320, and the encrypted ciphertext stored. Thus,
13 so long as the trusted core can retrieve gatekeeper storage key 320, it can decrypt
14 the hive keys 322, 324, and 326, and then use the hive keys to decrypt trusted
15 application secrets 328, 330, and 332.

16 Non-migrateable secrets 328 are unconditionally non-migrateable – they
17 cannot be transferred to another computing device. Non-migrateable secrets 328
18 are encrypted by an encryption algorithm that uses, as an encryption key, non-
19 migrateable key 322. The trusted core will not divulge non-migrateable key 322
20 to another computing device, so no other device will be able to decrypt trusted
21 application secrets 328. However, an upgraded trusted core (executing on the
22 same computing device) may still be able to access trusted application secrets 328
23 because, as discussed above, the upgraded trusted core will be able to retrieve

1 gatekeeper storage key 320. Although only a single non-migrateable key 322 is
2 illustrated, alternatively multiple non-migrateable keys may be used.

3 User-migrateable secrets 330 can be migrated/transferred to another
4 computing device, but only under the control or direction of the user. User-
5 migrateable key 324 can be transferred, under the control or direction of the user,
6 to another computing device. The encrypted trusted application secrets 330 can
7 also be transferred to the other computing device which, so long as the trusted core
8 of the other computing device has user-migrateable key 324, can decrypt trusted
9 application secrets 330.

10 Multiple user-migrateable keys 324 may be used. For example, each
11 trusted application that stores user-migrateable secrets may use a different user-
12 migrateable key (thereby allowing the migration of secrets for different trusted
13 applications to be controlled separately), or a single trusted application may use
14 different user-migrateable keys for different ones of its secrets. Which user-
15 migrateable key 324 to use to encrypt a particular trusted application secret is
16 identified by the trusted application when requesting secure storage of the secret.

17 In one implementation, this user control is created by use of a passphrase.
18 The user can input his or her own passphrase on the source computing device, or
19 alternatively the trusted core executing on the source computing device may
20 generate a passphrase and provide it to the user. The trusted core encrypts user-
21 migrateable key 324 to the passphrase, using the passphrase as the encryption key.
22 The ciphertext that is the encrypted trusted application secrets 330 can be
23 transferred to the destination computing device in any of a variety of manners
24 (e.g., copied onto a removable storage medium (e.g., optical or magnetic disk) and

1 the medium moved to and inserted into the destination computing device, copied
2 via a network connection, etc.).

3 The user also inputs the passphrase (regardless of who/what created the
4 passphrase) into the destination computing device. The encrypted user-
5 migrateable key 324 can then be decrypted by the trusted core at the destination
6 computing device using the passphrase. The trusted core at the destination device
7 can then encrypt user-migrateable key 324 using the gatekeeper storage key of the
8 trusted core at the destination device. Given user-migrateable key 324, the trusted
9 core at the destination device is able to retrieve secrets securely stored using key
10 324, assuming that the trusted core executing on the destination device is not a
11 different trusted core (or an earlier version of the trusted core) executing on the
12 source device. The retrieval of secrets is based on a manifest, as discussed in
13 more detail below.

14 The trusted core also typically authenticates the destination computing
15 device before allowing the encrypted user-migrateable key 324 to be transferred to
16 the destination computing device. Alternatively, at the user's discretion,
17 authentication of the destination computing device may not be performed. The
18 trusted core may perform the authentication itself, or alternatively rely on another
19 party (e.g., a remote authentication party trusted by the trusted core) to perform the
20 authentication or assist in the authentication.

21 The destination computing device can be authenticated in a variety of
22 different manners. In one implementation, the quote and/or pk_unseal operations
23 are used to verify that the trusted core executing on the destination computing
24 device is the same as or is known to the trusted core executing on the source
25

1 computing device (e.g., identified as or determined to be trustworthy to the trusted
2 core on the source computing device). The authentication may also involve
3 checking a list of “untrustworthy” certificates (e.g., a revocation list) to verify that
4 the trusted core on the destination computing device (based on its certificate) has
5 not been identified as being untrustworthy (e.g., broken by a mischievous user).
6 The authentication may also optionally include, analogous to verifying the
7 trustworthiness of the trusted core on the destination computing device, verifying
8 the trustworthiness of the destination computing device hardware (e.g., based on a
9 certificate of the hardware or platform), as well as verifying the trustworthiness of
10 one or more trusted applications executing on the destination computing device.

11 Third party-migrateable secrets 332 can be migrated/transferred to another
12 computing device, but only under the control or direction of a third party. This
13 third party could be the party that provided the secret to the trusted application, or
14 alternatively could be another party (such as a party that agrees to operate as a
15 controller/manager of how data is migrated amongst devices). Examples of third
16 party control include keys that control access to premium content (e.g., movies)
17 etc., which may be licensed to several of a user’s devices, and yet not freely
18 movable to any other device, or credentials used to log on to a corporate LAN
19 (Local Area Network), which can be moved, but only under the control of the
20 LAN administrator. This third party could also be another device, such as a
21 smartcard that tracks and limits the number of times the secret is migrated. Third
22 party-migrateable key 326 can be transferred, under the control or direction of the
23 third party, to another computing device. The encrypted trusted application
24 secrets 332 can also be transferred to the other computing device which, so long as

1 the trusted core of the other computing device has third party-migrateable key 326,
2 can decrypt trusted application secrets 332 (assuming that the trusted core
3 executing on the destination device is not a different trusted core (or an earlier
4 version of the trusted core) executing on the source device).

5 In one implementation, this user control is created by use of a public-
6 private key pair associated with the third party responsible for controlling
7 migration of secrets amongst machines. Multiple such third parties may exist,
8 each having its own public-private key pair and each having its own corresponding
9 third party-migrateable key 326. Each third party-migrateable key 326 has a
10 corresponding certificate 334 that includes the public key of the corresponding
11 third party. Each time that a trusted application requests secure storage of a third
12 party-migrateable secret, the trusted application identifies the third party that is
13 responsible for controlling migration of the secret. If a key 326 already exists for
14 the identified third party, then that key is used to encrypt the secret. However, if
15 no such key already exists, then a new key corresponding to the identified third
16 party is generated, added as one of keys 326, and is used to encrypt the secret.

17 In order to migrate a third party-migrateable secret, the trusted core
18 encrypts the third party-migrateable key 326 used to encrypt that secret with the
19 public key of the certificate 334 corresponding to the key 326. The ciphertext that
20 is the encrypted trusted application secrets 332 can be transferred to the
21 destination computing device in any of a variety of manners (e.g., copied onto a
22 removable storage medium (e.g., optical or magnetic disk) and the medium moved
23 to and inserted into the destination computing device, copied via a network
24 connection, etc.). The encrypted third party-migrateable key 326 is also

1 transferred to the destination computing device, and may be transferred along with
2 (or alternatively separately from) the encrypted trusted application secrets 332.

3 The trusted core executing on the source computing device, or alternatively
4 the third party corresponding to the encrypted third party-migrateable key, also
5 typically authenticates the destination computing device before allowing the
6 encrypted third party-migrateable key 326 to be transferred to the destination
7 computing device. Alternatively, at the discretion of the third party corresponding
8 to the encrypted third party-migrateable key, authentication of the destination
9 computing device may not be performed. The trusted core (or third party) may
10 perform the authentication itself, or alternatively rely on another party (e.g., a
11 remote authentication party trusted by the trusted core or third party) to perform or
12 assist in performing the authentication.

13 The trusted core executing on the destination computing device can then
14 access the third party corresponding to the encrypted third party-migrateable key
15 326 in order to have the key 326 decrypted. The third party can impose whatever
16 type of verification or other constraints that it desires in determining whether to
17 decrypt the key 326. For example, the third party may require the trusted core
18 executing on the destination computing device to authenticate itself, or may
19 decrypt the key 326 only if fewer than an upper limit number of computing
20 devices have requested to decrypt the key 326, or may require the user to verify
21 certain information over the telephone, etc.

22 If the third party refuses to decrypt the key 326, then the destination
23 computing device is not able to decrypt encrypted trusted application secrets 332.
24 However, if the third party does decrypt the key 326, then the third party returns
25

1 the decrypted key to the destination computing device. The decrypted key can be
2 returned in a variety of different secure methods, such as via a voice telephone call
3 between the user of the destination computing device and a representative of the
4 third party, using network security protocols (such as HTTPS (Secure HyperText
5 Transfer Protocol)), encrypting the key with a public key of a public-private key
6 pair of the destination computing device, etc. The trusted core at the destination
7 device can then encrypt third party-migrateable key 326 using the gatekeeper
8 storage key of the trusted core at the destination device.

9 Storing application secrets based on classes or types facilitates the
10 migration of the application secrets to other computing devices. Rather than using
11 a separate key for each application secret, the application secrets are classed
12 together, with only one key typically being needed for the user-migrateable class
13 and only one key per third party typically being needed for the third party-
14 migrateable class. Thus, for example, rather than requiring each user-migrateable
15 secret to have its own key that needs to be transferred to the destination device in
16 order to migrate the secrets to the destination device, only the single user-
17 migrateable key need be transferred to the destination device. Additionally, an
18 “all” class can also exist (e.g., associated with gatekeeper storage key 320 of Fig.
19 9) that allows all of the secrets (except the non-migrateable secrets) to be migrated
20 to the destination device by transferring and having decrypted only the gatekeeper
21 storage key (which can in turn be used to decrypt the encrypted hive keys). The
22 non-migrateable secrets can be kept from being migrated by not allowing the
23 encrypted non-migrateable hive key to be copied.

1 Fig. 10 illustrates an exemplary process 360 for securely storing secrets
2 using hive keys. The process of Fig. 10 is carried out by the trusted core of a
3 client computing device, and may be performed in software.

4 The first time the trusted core is booted, a gatekeeper storage key is
5 generated (act 362) and sealed, using a cryptographic measure of the trusted core,
6 to the trusted core (act 364). Eventually, a request to store a secret is received by
7 the trusted core from a trusted application (act 366), and the request includes an
8 identification of the type of secret (non-migrateable, user-migrateable, or third
9 party-migrateable). The trusted core generates a hive key for that type of secret if
10 needed (act 368). A hive key is needed if no hive key of that type has been
11 created by the trusted core yet, or if the identified user-migrateable key has not
12 been created yet, or if a hive key corresponding to the third party of a third party-
13 migrateable secret has not been created yet.

14 Once the correct hive key is available, the trusted core uses the hive key to
15 encrypt the trusted application secret (act 370). Additionally, the trusted core uses
16 the gatekeeper storage key to encrypt the hive key (act 372).

17 Fig. 11 illustrates an exemplary process 400 for migrating secrets from a
18 source computing device to a destination computing device. The process of Fig.
19 11 is carried out by the trusted cores on the two computing devices. The process
20 of Fig. 11 is discussed with reference to components of Fig. 9.

21 Initially, a request to migrate or transfer secrets to a destination computing
22 device is received at the source computing device (act 402). The trusted core on
23 the source computing device determines whether/how to allow the transfer of
24 secrets based on the type of secret (act 404). If the secret is a non-migrateable
25

1 secret, then the trusted core does not allow the secret to be transferred or migrated
2 (act 406).

3 If the secret is a user-migrateable secret, then the trusted core obtains a user
4 passphrase (act 408) and encrypts the hive key corresponding to the secret using
5 the passphrase (act 410). The trusted core also authenticates the destination
6 computing device as being trusted to receive the secret (act 412). If the
7 destination computing device is not authenticated, then the trusted core does not
8 transfer the encrypted hive key to the destination computing device. Assuming the
9 destination computing device is authenticated, the encrypted hive key as well as
10 the encrypted secret is received at the destination computing device (act 414), and
11 the trusted core at the destination computing device also receives the passphrase
12 from the user (act 416). The trusted core at the destination computing devices
13 uses the passphrase to decrypt the hive key (act 418), thereby allowing the trusted
14 core to decrypt the encrypted secrets when requested.

15 If the secret is a third party-migrateable secret, then the trusted core on the
16 source computing device encrypts the hive key corresponding to the secret using
17 the public key of the corresponding third party (act 420). The trusted core on the
18 source computing device, or alternatively the third party corresponding to the hive
19 key, also authenticates the destination computing device (act 422). If the
20 destination computing device is not authentication then the trusted core does not
21 transfer the encrypted hive key to the destination computing device (or
22 alternatively, the third party does not decrypt the hive key). Assuming the
23 destination computing device is authenticated, the encrypted hive key as well as
24 the encrypted secret is received at the destination computing device (act 424). The

1 trusted core at the destination computing device contacts the corresponding third
2 party to decrypt the hive key (act 426), thereby allowing the trusted core to
3 decrypt the encrypted secrets when requested.

4 Thus, secure secret storage is maintained by allowing trusted processes to
5 restrict whether and how trusted application secrets can be migrated to other
6 computing devices, and by the trusted core enforcing such restrictions.
7 Furthermore, migration of secrets to other computing devices is facilitated by the
8 use of the gatekeeper storage key and hive keys, as only one or a few keys need to
9 be moved in order to have access to the application secrets held by the source
10 device. Additionally, the use of hive keys to migrate secrets to other computing
11 devices does not interfere with the ability of the trusted applications or the trusted
12 core to authenticate itself to third parties.

13

14 **Backup of Secrets**

15 Situations can arise where the hardware or software of a client computing
16 device 102 of Fig. 1 is damaged or fails. Because of the possibility of such
17 situations arising, it is generally prudent to backup the data stored on client
18 computing device 102, including the securely stored secrets. However, care
19 should be taken to ensure that the backup of the securely stored secrets does not
20 compromise the security of the storage.

21 There are two primary situations that data backups are used to recover
22 from. The first is the failure of the mass storage device that stores the trusted core
23 (e.g., a hard disk) or the operating system executing on the computing device, and
24 the second is the damaging of the device sufficiently to justify replacement of the

1 computing device with a new computing device (e.g., a heavy object fell on the
2 computing device, or a power surge destroyed one or more components).

3 In order to recover from the first situation (failure of the mass storage
4 device or operating system), the contents of the mass storage device (particularly
5 the trusted core and the trusted application secrets) are backed up when the
6 computing device is functioning properly. Upon failure of the mass storage device
7 or operating system, the mass storage device can be erased (e.g., formatted) or
8 replaced, and the backed up data stored to the newly erased (or new) mass storage
9 device. Alternatively, rather than backing up the trusted core, the computing
10 device may have an associated "recovery" disk (or other media) that the
11 manufacturer provides and that can be used to copy the trusted core from when
12 recovering from a failure. When the computing device is booted with the backed
13 up data, the trusted core will have the same digest as the trusted core prior to the
14 failure, so that the new trusted core will be able to decrypt the gatekeeper storage
15 key and thus the trusted application secrets.

16 In order to recover from the second situation (replacement of the
17 computer), the backing up of securely stored secrets is accomplished in a manner
18 very similar to the migration of secrets from one computing device to another. In
19 the situation where the computing device 102 is damaged and replaced with
20 another computing device, the backing up is essentially migrating the trusted
21 application secrets from a source computing device (the old, damaged device) to a
22 destination computing device (the new, replacement device).

23 Recovery from the second situation varies for different trusted application
24 secrets based on the secret types. Non-migrateable secrets are not backed up.
25

1 This can be accomplished by the trusted core not allowing the non-migrateable
2 secrets to be copied from the computing device, or not allowing the non-
3 migrateable key to be copied from the computing device, when backing up data.

4 User-migrateable secrets are backed up using a passphrase. During the
5 backup procedure, a user passphrase(s) is obtained and used to encrypt the user-
6 migrateable key(s), with the encrypted keys being stored on a backup medium
7 (e.g., a removable storage medium such as a disk or tape, a remote device such as
8 a file server, etc.). To recover the backup data, the user can copy the backed up
9 encrypted trusted application secrets, as well as the user-migrateable key(s)
10 encrypted to the passphrase(s), to any other device he or she desires. Then, by
11 entering the passphrase(s) to the other device, the user can allow the trusted core
12 to decrypt and retrieve the trusted application secrets.

13 Third party-migrateable secrets are backed up using a public key(s) of the
14 third party or parties responsible for controlling the migration of the secrets.
15 During the backup procedure, the trusted core encrypts the third party-migrateable
16 key(s) with the public key(s) of the corresponding third parties, and the encrypted
17 keys are stored on a backup medium (e.g., a removable storage medium such as a
18 disk or tape, a remote device such as a file server, etc.). To recover the backup
19 data, the user can copy the backed up encrypted trusted application secrets to any
20 other device he or she desires, and contact the appropriate third party or parties to
21 decrypt the encrypted keys stored on the backup medium. Assuming the third
22 party or parties authorize the retrieval of the keys, the third party or parties decrypt
23 the keys and return (typically in a secure manner) the third party-migrateable

1 key(s) to the other computing device, which the trusted core can use to decrypt
2 and retrieve the trusted application secrets.

3 Thus, analogous to the discussion of hive keys and secret migration above,
4 trusted processes are allowed to restrict whether and how trusted application
5 secrets can be backed up, and the trusted core enforces such restrictions.
6 Additionally, the backing up of secrets does not interfere with the ability of the
7 trusted applications or the trusted core to authenticate itself to third parties.

8

9 **Manifests and Application Security Policies**

10 Oftentimes, trusted application components and modules are more likely to
11 be upgraded than are components and modules of the trusted core. Trusted
12 applications frequently include various dynamic link libraries (DLL's), plug-ins,
13 etc. and allow for different software configurations, each of which can alter the
14 binaries which execute as the trusted application. Using a digest for the trusted
15 application can thus be burdensome as the digest would be changing every time
16 one of the binaries for the trusted application changes. Thus, rather than using a
17 digest for the trusted applications as is described above for the trusted core, a
18 security model is defined for trusted applications that relies on manifests. A
19 manifest is a policy statement which attempts to describe what types of binaries
20 are allowed to be loaded into a process space for a trusted application. This
21 process space is typically a virtual memory space, but alternatively may be a non-
22 virtual memory space. Generally, the manifest specifies a set of binaries, is
23 uniquely identifiable, and is used to gate access to secrets. Multiple manifests can
24 be used in a computing device at any one time – one manifest may correspond to
25

1 multiple different applications (sets of binaries), and one application (set of
2 binaries) may correspond to multiple different manifests.

3 Fig. 12 illustrates an exemplary manifest 450 corresponding to a trusted
4 application. Manifest 450 can be created by anybody - there need not be any
5 restrictions on who can create manifests. Certain trust models may insist on
6 authorization by some given authority in order to generate manifests. However,
7 this is not an inherent property of manifests, but a way of using them - in principle,
8 no authorization is needed to create a manifest. Manifest 450 includes several
9 portions: an identifier portion 452 made up of a triple (K, U, V), a signature
10 portion 454 including a digital signature over manifest 450 (except for signature
11 portion 454), a digest list portion 456, an export statement list portion 458, and a
12 set of properties portion 460. An entry point list 462 may optionally be included.

13 Identifier portion 452 is an identifier of the manifest. In the illustrated
14 example the manifest identifier is a triple (K, U, V), in which K is a public key of
15 a public-private key pair of the party that generates manifest 450. U is an arbitrary
16 identifier. Generally, U is a member of a set M_u , where the exact definition of M_u
17 is dependent upon the specific implementation. One condition on set M_u is that all
18 of its elements have a finite representation (that is, M_u is countable). M_u could be,
19 for example, the set of integers, the set of strings of finite length over the Latin
20 alphabet, the set of rational numbers, etc. In one implementation, the value U is a
21 friendly name or unique identifier of the party that generates manifest 450. V is
22 similar to U, and can be a member of a set M_v having the same conditions as M_u
23 (which may be the same set that U is a member of, or alternatively a different set).
24 Additionally, there is an (total or partial) defined on the set M_v (e.g., increasing

numerical order, alphabetical order, or some arbitrarily defined order). In one implementation, V is the version number of manifest 450. The trusted application corresponding to manifest 450 is identified by the triple in portion 452.

Manifest identifier portion 452 is described herein primarily with reference to the triple (K, U, V). Alternatively, manifest identifiers may not include all three elements K, U, and V. For example, if version management is not needed, the V component can be omitted.

Alternatively, different manifest identifiers may also be used. For example, any of a variety of conventional cryptographic hashing functions (such as SHA-1) may be used to generate a hash of one or more portions of manifest 450 (e.g., portion 456). The resultant hash value can be used as the manifest identifier.

Signature portion 454 includes a digital signature over the portions of manifest 450 other than signature portion 454 (that is, portions 452, 456, 458, and 460). Alternatively, one or more other portions of manifest 450 may also be excluded from being covered by the digital signature, such as portion 458. The digital signature is generated by the party that generates manifest 450, and is generated using the private key corresponding to the public key K in portion 452. Thus, given manifest 450, a device (such as a trusted core) can verify manifest 450 by checking the manifest signature 454 using the public key K. Additionally, this verification may be indirectioned through a certificate chain.

Alternatively, a digital signature over a portion(s) of manifest 450 may not be included in manifest 450. The digital signature in portion 454 serves to tie lists portion 456 to the manifest identifier. In various alternatives, other mechanisms may be used to tie lists portion 456 to the manifest identifier. For example, if the

1 manifest identifier is a hash value generated by hashing portion 456, then the
2 manifest identifier inherently ties lists portion 456 to the manifest identifier.

3 Certificate lists 456 are two lists (referred to as S and T) of public key
4 representations. In one implementation, lists 456 are each a list of certificate
5 hashes. The S list is referred to as an inclusion list while the T list is referred to as
6 an exclusion list. The certificate hashes are generated using any of a wide variety
7 of conventional cryptographic hashing operations, such as SHA-1. List S is a list
8 of hashes of certificates that certify the public key which corresponds to the
9 private key that was used to sign the certificates in the chain that corresponds to
10 the binaries that are authorized by manifest 450 to execute in the virtual memory
11 space. A particular manufacturer (e.g., Microsoft Corporation) may digitally sign
12 multiple binaries using the same private key, and thus the single certificate that
13 includes the public key corresponding to this private key may be used to authorize
14 multiple binaries to execute in the virtual memory space. Alternatively, a
15 manufacturer can generate an entirely new key for each binary which is
16 subsequently deleted. This will result in the same mechanism being used to
17 identify a single, unique application as opposed to one from a family. The “hash-
18 of-a-certificate” scheme is hence a very flexible scheme for describing
19 applications or families of applications.

20 List T is a list of hashes of certificates that certify the public key which
21 corresponds to the private key that was used to sign the certificates in the chain
22 that corresponds to the binaries that are not authorized by manifest 450 to execute
23 in the virtual memory space. List T may also be referred to as a revocation list.
24 Adding a particular certificate to list T thus allows manifest 450 to particularly
25

1 identify one or more binaries that are not allowed to execute in the virtual memory
2 space. The entries in list T override the entries in list S. Thus, in order for a
3 binary to be authorized to execute in a virtual memory space corresponding to
4 manifest 450, the binary must have a certificate hash that is the same as a
5 certificate hash in list S (or have a certificate that identifies a chain of one or more
6 additional certificates, at least one of which is in list S) but is not the same as any
7 certificate hash in list T. In addition, none of the certificates in the chain from the
8 certificate in S to the leaf certificate that contains the hash of the binary can be
9 contained in list T. If both of these conditions are not satisfied, then the binary is
10 not authorized to execute in the virtual memory space corresponding to manifest
11 450.

12 The T list, in conjunction with the S list, can be used flexibly. For
13 example, given an inclusion of all applications certified by a given root in the
14 inclusion list (S), the exclusion list (T) can be used to exclude one or more
15 applications that are known to be vulnerable or have other bugs. Similarly, given
16 a certification hierarchy, with the root certificate on the inclusion list (S), the
17 exclusion list (T) can be used to remove one or more of the child keys in the
18 hierarchy (and binaries certified by them) that have been compromised.

19 Alternatively, other public key representations or encodings besides
20 certificate hashes can be used as one or both of the S and T lists. For example,
21 rather than certificate hashes, the S and T lists may be the actual certificates that
22 certify the public keys which correspond to the private keys that were used to sign
23 the certificates in the chains that correspond to the binaries that are authorized by
24 manifest 450 to execute (the S list) or not execute (the T list) in the virtual
25

1 memory space. By way of another example, the S and T lists may be just the
2 public keys which correspond to the private keys that were used to sign the
3 certificates in the chains that correspond to the binaries that are authorized by
4 manifest 450 to execute (the S list) or not execute (the T list) in the virtual
5 memory space.

6 Export statement list portion 458 includes a list of zero or more export
7 statements that allow a trusted application secret associated with manifest 450 to
8 be exported (migrated) to another trusted application on the same computing
9 device. Each trusted application executing on a client computing device 102 of
10 Fig. 1 has a corresponding manifest 450, and thus each trusted application secret
11 securely saved by the trusted application is associated with manifest 450. Export
12 statement list portion 458 allows the party that generates manifest 450 to identify
13 the other trusted applications to which the trusted application secrets associated
14 with manifest 450 can be exported and made available for retrieving.

15 Each export statement includes a triple (A, B, S), where A is the identifier
16 (K, U, V) of the source manifest, B is the identifier (K, U, V) of the destination
17 manifest, and S is a digital signature over the source and destination manifest
18 identifiers. B may identify a single destination manifest, or alternatively a set of
19 destination manifests. Additionally, for each (K, U) in B, a (possibly open)
20 interval of V values may optionally be allowed (e.g., “version 3 and higher”, or
21 “versions 2 through 5”). The digital signature S is made using the same private
22 key as was used to sign manifest 450 (in order to generate the signature in portion
23 454).

1 Export statements may be device-independent and thus not limited to being
2 used on any particular computing device. Alternatively, an export statement may
3 be device-specific, with the export statement being useable on only one particular
4 computing device (or set of computing devices). This one particular computing
5 device may be identified in different manners, such as via a hardware id or a
6 cryptographic mechanism (e.g., the export statement may be encrypted using the
7 public key associated with the particular computing device). If a hardware id is
8 used to identify a particular computing device, the export statement includes an
9 additional field which states the hardware id (thus, the issuer of the manifest could
10 control on a very fine granularity who can move secrets).

11 Additionally, although illustrated as part of manifest 450, one or more
12 export statements may be separate from, but associated with, manifest 450. For
13 example, the party that generates manifest 450 may generate one or more export
14 statements after manifest 450 is generated and distributed. These export
15 statements are associated with the manifest 450 and thus have the same affect as if
16 they were included in manifest 450. For example, a new trusted application may
17 be developed after the manifest 450 is generated, but the issuer of the manifest 450
18 would like the new trusted application to have access to secrets from the
19 application associated with the manifest 450. The issuer of the manifest 450 can
20 then distribute an export statement (e.g., along with the new trusted application or
21 alternatively separately) allowing the secrets to be migrated to the new trusted
22 application.

23 If a user or trusted application desires to export trusted application secrets
24 from a source trusted application to a destination trusted application, the trusted
25

1 core checks to ensure that the manifest identifier of the desired destination trusted
2 application is included in export statement list portion 758. If the manifest
3 identifier of the desired destination trusted application is included in export
4 statement list portion 758, then the trusted core allows the destination trusted
5 application to have access to the source trusted application secrets; otherwise, the
6 trusted core does not allow the destination trusted application to have access to the
7 source trusted application secrets. Thus, although a user may request that trusted
8 application secrets be exported to another trusted application, the party that
9 generates the manifest for the trusted application has control over whether the
10 secrets can actually be exported to the other trusted application.

11 Properties portion 460 identifies a set of zero or more properties for the
12 manifest 450 and/or executing process corresponding to manifest 450. Various
13 properties can be included in portion 460. Example properties include: whether
14 the process is debuggable, whether to allow (or under what conditions to allow)
15 additional binaries to be added to the virtual memory space after the process
16 begins executing, whether to allow implicit upgrades to higher manifest version
17 numbers (e.g., allow upgrades from one manifest to another based on the K and U
18 values of identifier 452, without regard for the V value), whether other processes
19 (and what other processes) should have access to the virtual memory space of the
20 process (e.g. to support secure shared memory), what/whether other resources
21 should be shareable (e.g. “pipe” connections, mutexes (mutually exclusives), or
22 other OS resources), and so forth.

23 Entry point list 462 is optional and need not be included in manifest 450.
24 In one implementation, an entry point list is included in the binary or a certificate
25

for the binary, and thus not included in manifest 450. However, in alternative embodiments entry point list 462 may be included in manifest 450. Entry point list 462 is a list of entry points into the executing process. Entry point list 462 is typically generated by the party that generates manifest 450. These entry points can be stored in a variety of different manners, such as particular addresses (e.g., offsets relative to some starting location, such as the beginning address of a particular binary), names of functions or procedures, and so forth. These entry points are the only points of the process that can be accessed by other processes (e.g., to invoke functions or methods of the process). When a request to access a particular address in the virtual memory space of an executing process associated with manifest 450 is received, the trusted core checks whether the particular address corresponds to an entry point in entry point list 462. If the particular address does correspond to an entry point in entry point list 462, then the access is allowed; otherwise, the trusted core denies the access.

The manifest is used by the trusted core in controlling authentication of trusted application processes and access to securely stored secrets by trusted application processes executing on the client computing device. When referencing a trusted application process, the trusted core (or any other entity) can refer to its identifier (the triple K, U, V). The trusted core exposes versions of the Seal, Unseal, Quote, and Unwrap operations analogous to those primitive operations discussed above, except that it is the trusted core that is exposing the operations rather than the underlying hardware of the computing device, and the parameters may vary. In one implementation, the versions of the Seal, Unseal, Quote, and

1 Unwrap operations that are exposed by the trusted core and that can be invoked by
2 the trusted application processes are as follows.

3 The Seal operation exposed by the trusted core takes the following form:

4 *Seal (secret, public_key (K), identifier, version, secret_type)*

5 where *secret* represents the secret to be securely stored, *public_key (K)* represents
6 the K component of a manifest identifier, *identifier* represents the U component of
7 a manifest identifier, *version* represents the V value of a manifest identifier, and
8 *secret_type* is the type of secret to be stored (e.g., non-migrateable, user-
9 migrateable, or third party-migrateable). The manifest identifier (the K, U, and V
10 components) is a manifest identifier as described above (e.g., with reference to
11 manifest 450). The K and U portions of the manifest identifier refer to the party
12 that generated the manifest for the process storing the secret, while the V portion
13 refers to the versions of the manifest that should be allowed to retrieve the secret.
14 In the general case, the (K,U,V) triple can be a list of such triples and the value V
15 can be a range of values.

16 When the Seal operation is invoked, the trusted core encrypts the secret and
17 optionally additional parameters of the operation using the appropriate hive key
18 based on the *secret_type*. The encrypted secret is then stored by the trusted core in
19 secret store 126 of Fig. 2 or 146 of Fig. 3 cryptographically bound to the
20 associated rules (the list {(K,U,V)}), or alternatively in some other location.

21 The Unseal operation exposed by the trusted core takes the following form:

22 *Unseal (encrypted_secret)*

23 where *encrypted_secret* represents the ciphertext that has encrypted in it the secret
24 to be retrieved together with the (K, U, V) list that names the application(s)

1 qualified to retrieve the secret. In response to the Unseal operation, the trusted
2 core obtains the encrypted secret and determines whether to reveal the secret to the
3 requesting process. The trusted core reveals the secret to the requesting process
4 under two different sets of conditions; if neither of these sets of conditions is
5 satisfied then the trusted core does not reveal the secret to the requesting process.
6 The first set of conditions is that the requesting process was initiated with a
7 manifest that is properly formed and is included in the (K, U, V) list (or the K, U,
8 V value) indicated by the sealer. This is the common case: An application can
9 seal a secret naming its own manifest, or all possible future manifests from the
10 same software vendor. In this case, the same application or any future application
11 in the family has automatic access to its secrets.

12 The second set of conditions allows a manifest issuer to make a specific
13 allowance for other applications to have access to the secrets previously sealed
14 with more restrictive conditions. This is managed by an export certificate, which
15 provides an override that allows secrets to be migrated to other applications from
16 other publishers not originally named in the (K, U, V) list of the sealer. To avoid
17 uncontrolled and insecure migration, export lists should originate from the
18 publisher of the original manifest. This restriction is enforced by requiring that the
19 publisher sign the export certificate with the key originally used to sign the
20 manifest of the source application. This signature requirement may also be
21 indirectioned through certificate chains.

22 To process an export certificate, the trusted core is a) furnished with the
23 manifest from the original publisher (i.e., the manifest issuer), b) furnished with
24 the export certificate itself which is signed by the original publisher, and c)
25

1 running a process that is deemed trustworthy in the export certificate. If all these
2 requirements are met, the running process has access to the secrets sealed by the
3 original process.

4 The Quote and Unwrap operations provide a way for the trusted core to
5 authenticate to a third party that it is executing a trusted application process with a
6 manifest that meets certain requirements.

7 The Unwrap operation uses ciphertext as its single parameter. A third
8 (arbitrary) party initially generates a structure that includes five parts: a *secret*, a
9 *public_key K*, an *identifier U*, a *version V*, and a *hive_id*. Here, *secret* represents
10 the secret to be revealed if the appropriate conditions are satisfied, *public_key K*
11 represents the public key of the party that needs to have digitally signed the
12 manifest for the process, *identifier U* is the identifier of the party that needs to
13 have generated the manifest for the process, *version V* is a set of zero or more
14 acceptable versions of the manifest, and *hive_id* is the type of secret being
15 revealed (e.g., non-migrateable, user-migrateable, or third party-migrateable). The
16 party then encrypts this structure using the public key of the public-private key
17 pair known to belong to a trustworthy trusted core (presumably because of
18 certification of the public part of this key). The manner in which the trusted core
19 gets this key is discussed in additional detail in U.S. Patent Application No.
20 09/227,611 entitled "Loading and Identifying a Digital Rights Management
21 Operating System" and U.S. Patent Application No. 09/227,561 entitled "Digital
22 Rights Management Operating System". A trusted application receives the
23 ciphertext generated by the third party and invokes the Unwrap operation exposed
24 by the trusted core.

1 The trusted core responds to the Unwrap operation by using its private key
2 of the public-private key pair to decrypt the ciphertext received from the invoking
3 party. The trusted core compares the conditions in or associated with the
4 encrypted ciphertext to the manifest associated with the appropriate trusted
5 application process. The appropriate trusted application process can be identified
6 explicitly by the third party that generated the ciphertext being unwrapped, or
7 alternatively inherently as the trusted application invoking the Unwrap operation
8 (so the trusted core knows that whichever process invokes the Unwrap operation is
9 the appropriate trusted application process). If the manifest associated with the
10 process satisfies all of the conditions in the encrypted ciphertext, then the process
11 is authorized to retrieve the secret, and the trusted core provides the secret to the
12 process. However, if one or more of the conditions in the encrypted ciphertext are
13 not satisfied by the manifest associated with the process, then the process is not
14 authorized to retrieve the secret and the trusted core does not provide the secret to
15 the process.

16 In addition to manifest-based conditions, the Unwrap operation may also
17 have conditions on the data of the secret. If the conditions on the data (e.g., to
18 verify its integrity) are not satisfied then the trusted core does not provide the
19 secret to the process (even if the manifest conditions are satisfied). For example,
20 the encrypted secret may include both the data of the secret and a cryptographic
21 hash of the data. The trusted core verifies the integrity of the data by hashing the
22 data and verifying the resultant hash value.

23 The Unwrap operation naming the manifest or manifests of the
24 application(s) allowed to decrypt the secret allows a remote party to conveniently
25

1 express that a secret should only be revealed to a certain application or set of
2 applications on a particular host computer running a particular trusted core.

3 An alternative technique is based on the use of the quote operation, which
4 allows an application value to be cryptographically associated with the manifest of
5 the application requesting the quote operation. The quote operation associates an
6 application-supplied value with an identifier for the running software. When
7 previously introduced, the quote operation was implemented in hardware, and
8 allowed the digest of the trusted core to be cryptographically associated with some
9 trusted core-supplied data. When implemented by the trusted core on behalf of
10 applications, the quote operation will generate a signed statement that a particular
11 value X was supplied by a process running under a particular manifest (K, U, V),
12 where the value X is an input parameter to the quote operation. The value X can
13 be used as part of a more general authentication protocol. For example, such a
14 statement can be sent as part of a cryptographic interchange between a client and a
15 server to allow the server to determine that the client it is talking to is a good
16 device running a trusted core, and an application that it trusts before revealing any
17 secret data to it. The requesting party can analyze the manifest and make its own
18 determination of whether it is willing to trust the process.

19 Fig. 13 illustrates an exemplary process 500 for controlling execution of
20 processes in an address space based on a manifest. The process of Fig. 13 is
21 discussed with reference to components in Fig. 12, and is implemented by a
22 trusted core.

23 Initially, a request to execute a process is received by the trusted core (act
24 502). This request may be received from a user or alternatively another process

executing on the same client computing device as the trusted core or alternatively on another computing device in communication with the client computing device. In response to the request, a virtual memory space for the process is set up by the trusted core (act 504) and the binaries necessary to execute the process are loaded into the virtual memory space (act 506). It should be noted that, in act 506, the binaries are loaded into the memory space but execution of the binaries has not yet begun. The trusted core then initializes the environment and obtains a manifest for the process (act 508). Typically, the manifest is provided to the trusted core as part of the request to execute the process.

The trusted core checks whether all of the loaded binaries are consistent with the manifest (act 510). In one implementation, this check for consistency involves verifying that the certificate (or certificate hash) of each binary is in the S list in portion 456 of manifest 450, and that certificates (or certificate hashes) for none of the binaries are in the T list in portion 456. This certificate verification may be indirected through a certificate list. If the loaded binaries are not consistent with the manifest (e.g., at least one is not in the S list and/or at least one is in the T list), then process 500 fails – the requested process is not executed (act 512).

However, if the loaded binaries are consistent with the manifest, then the trusted core allows the processor to execute the binaries in the virtual memory space (act 514). Execution of the loaded binaries typically is triggered by an explicit request from an outside entity (e.g. another process). A request may be subsequently received, typically from the executing process or some other process, to load an additional binary into the virtual memory space. The trusted core

1 continues executing the process if no such request is received (acts 514 and 516).
2 However, when such a request is received, the trusted core checks whether the
3 additional binary is consistent with manifest 450 (act 518). Consistency in act 518
4 is determined in the same manner as act 510 – the additional binary is consistent
5 with manifest 450 if its certificate (or certificate hash) is in the S list in portion 456
6 of manifest 450 and is not in the T list in portion 456.

7 If the additional binary is not consistent with manifest 450, then the
8 additional binary is not loaded into the virtual memory space and allowed to
9 execute, and processing continues to act 514. However, if the additional binary is
10 consistent with manifest 450, then the additional binary is loaded into the virtual
11 memory space (act 520), and processing of the binaries (including the additional
12 binary) continues.

13 Alternatively, rather than loading the binaries (act 506) and checking
14 whether the loaded binaries are consistent with the manifest (act 510), the manifest
15 can be obtained prior to loading the binaries into the virtual memory space (e.g.,
16 provided as part of the initial request to execute a trusted process in act 502). In
17 this case, each request to load a binary is checked against the manifest. Binaries
18 which are not allowed by the manifest are not loaded into the virtual memory
19 space, whereas binaries that are allowed are loaded into the virtual memory space.

20 Fig. 14 illustrates an exemplary process 540 for upgrading to a new version
21 of a trusted application. The process of Fig. 14 is discussed with reference to
22 components in Fig. 12, and is implemented by a computing device (typically other
23 than the client computing device). Typically, the upgraded version of a trusted

1 application is prepared by the same party that prepared the previous version of the
2 trusted application.

3 Initially, a trusted application upgrade request is received along with one or
4 more new components or modules (e.g., binaries) for the trusted application to be
5 upgraded (act 542). These new components or modules may replace previous
6 versions of the components or modules in the previous version of the process, or
7 alternatively may be new components or modules that have no counterpart in the
8 previous version. A party begins generating a new manifest 450' for the new
9 version of the trusted application including a new triple (K', U', V') identifier for
10 the new version and appropriate certificate hashes (or alternatively certificates) in
11 the appropriate S and T lists in portion 456 (act 544). Oftentimes (e.g., when the
12 issuer of the new manifest is also the issuer of the old manifest and chooses K=K')
13 the K' and U' parts of the triple will be the same as the K and U parts of the triple
14 identifier of the previous version, so that only V and V' differ (that is, only the
15 versions in the identifier differ). The new manifest 450' is then made available to
16 the client computing device(s) where the new version of the trusted application is
17 to be executed (act 546).

18 Generally, there are three situations for application upgrades. The first
19 situation is where some binaries for the application are changed, added, and/or
20 removed, but the old manifest allows the new binaries to be loaded and loading the
21 old binaries is not considered to harm security. In this situation, the manifest does
22 not have to change at all and no secrets have to be migrated. The user simply
23 installs the new binaries on his machine and they are allowed to execute.

1 The second situation is where some binaries are changed, added, and/or
2 removed, and the old manifest is no longer acceptable because some of the old
3 binaries (which can still be loaded under the old manifest) compromise security
4 and/or some of the changed or new binaries cannot be loaded under the old
5 manifest. The issuer of the old manifest decides to issue a new manifest with the
6 same K,U. Initially, the software manufacturer produces new binaries. These new
7 binaries are digitally signed (certificates are issued) and a new manifest is created.
8 This new manifest (via its S and T lists) allows the new binaries to be executed but
9 does not allow the old binaries to be executed (at least not the binaries that
10 compromise security). It should be noted that there is no inherent relationship
11 between the S and T lists of the old manifest and the S and T lists of the new
12 manifest. It should also be noted that, if the S list is completely changed in the
13 new manifest, and some old binaries are re-used, the old binaries may need to be
14 signed with a new private key.

15 A user then receives all three things (the new binaries, the certificates for
16 the new binaries, and the new manifest) and installs all three on his or her
17 machine. Secrets do not have to be migrated, because the new manifest is just a
18 new version of the old one. The new binaries are allowed to execute, but the old
19 binaries are not.

20 The third situation is where secrets have to be migrated between different
21 applications that are not versions of each other. This situation is handled as
22 described above regarding export statements.

23 Thus, secure secret storage is maintained by the trusted core imposing
24 restrictions, based on the manifests, on which trusted processes can retrieve
25

1 particular secrets. The manifests also provide a way for trusted applications to be
2 authenticated to remote parties.

3

4 **Exemplary Computing Device**

5 Fig. 15 illustrates a general exemplary computer environment 600, which
6 can be used to implement various devices and processes described herein. The
7 computer environment 600 is only one example of a computing environment and
8 is not intended to suggest any limitation as to the scope of use or functionality of
9 the computer and network architectures. Neither should the computer
10 environment 600 be interpreted as having any dependency or requirement relating
11 to any one or combination of components illustrated in the exemplary computer
12 environment 600.

13 Computer environment 600 includes a general-purpose computing device in
14 the form of a computer 602. Computer 602 can be, for example, a client
15 computing device 102 or server device 104 of Fig. 1, a device used to generate a
16 trusted application or manifest, etc. The components of computer 602 can include,
17 but are not limited to, one or more processors or processing units 604, a system
18 memory 606, and a system bus 608 that couples various system components
19 including the processor 604 to the system memory 606.

20 The system bus 608 represents one or more of any of several types of bus
21 structures, including a memory bus or memory controller, a peripheral bus, an
22 accelerated graphics port, and a processor or local bus using any of a variety of
23 bus architectures. By way of example, such architectures can include an Industry
24 Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an

1 Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA)
2 local bus, and a Peripheral Component Interconnects (PCI) bus also known as a
3 Mezzanine bus.

4 Computer 602 typically includes a variety of computer readable media.
5 Such media can be any available media that is accessible by computer 602 and
6 includes both volatile and non-volatile media, removable and non-removable
7 media.

8 The system memory 606 includes computer readable media in the form of
9 volatile memory, such as random access memory (RAM) 610, and/or non-volatile
10 memory, such as read only memory (ROM) 612. A basic input/output system
11 (BIOS) 614, containing the basic routines that help to transfer information
12 between elements within computer 602, such as during start-up, is stored in ROM
13 612. RAM 610 typically contains data and/or program modules that are
14 immediately accessible to and/or presently operated on by the processing unit 604.

15 Computer 602 may also include other removable/non-removable,
16 volatile/non-volatile computer storage media. By way of example, Fig. 15
17 illustrates a hard disk drive 616 for reading from and writing to a non-removable,
18 non-volatile magnetic media (not shown), a magnetic disk drive 618 for reading
19 from and writing to a removable, non-volatile magnetic disk 620 (e.g., a “floppy
20 disk”), and an optical disc drive 622 for reading from and/or writing to a
21 removable, non-volatile optical disc 624 such as a CD-ROM, DVD-ROM, or other
22 optical media. The hard disk drive 616, magnetic disk drive 618, and optical disc
23 drive 622 are each connected to the system bus 608 by one or more data media
24 interfaces 626. Alternatively, the hard disk drive 616, magnetic disk drive 618,
25

1 and optical disc drive 622 can be connected to the system bus 608 by one or more
2 interfaces (not shown).

3 The various drives and their associated computer storage media provide
4 non-volatile storage of computer readable instructions, data structures, program
5 modules, and other data for computer 602. Although the example illustrates a
6 hard disk 616, a removable magnetic disk 620, and a removable optical disc 624, it
7 is to be appreciated that other types of computer readable media which can store
8 data that is accessible by a computer, such as magnetic cassettes or other magnetic
9 storage devices, flash memory cards, CD-ROM, digital versatile discs (DVD) or
10 other optical storage, random access memories (RAM), read only memories
11 (ROM), electrically erasable programmable read-only memory (EEPROM), and
12 the like, can also be utilized to implement the exemplary computing system and
13 environment.

14 Any number of program modules can be stored on the hard disk 616,
15 magnetic disk 620, optical disc 624, ROM 612, and/or RAM 610, including by
16 way of example, an operating system 626, one or more application programs 628
17 (e.g., trusted applications), other program modules 630, and program data 632.
18 Each of such operating system 626, one or more application programs 628, other
19 program modules 630, and program data 632 (or some combination thereof) may
20 implement all or part of the resident components that support the distributed file
21 system.

22 A user can enter commands and information into computer 602 via input
23 devices such as a keyboard 634 and a pointing device 636 (e.g., a "mouse").
24 Other input devices 638 (not shown specifically) may include a microphone,

1 joystick, game pad, satellite dish, serial port, scanner, and/or the like. These and
2 other input devices are connected to the processing unit 604 via input/output
3 interfaces 640 that are coupled to the system bus 608, but may be connected by
4 other interface and bus structures, such as a parallel port, game port, or a universal
5 serial bus (USB).

6 A monitor 642 or other type of display device can also be connected to the
7 system bus 608 via an interface, such as a video adapter 644. In addition to the
8 monitor 642, other output peripheral devices can include components such as
9 speakers (not shown) and a printer 646 which can be connected to computer 602
10 via the input/output interfaces 640.

11 Computer 602 can operate in a networked environment using logical
12 connections to one or more remote computers, such as a remote computing device
13 648. By way of example, the remote computing device 648 can be a personal
14 computer, portable computer, a server, a router, a network computer, a peer device
15 or other common network node, and the like. The remote computing device 648 is
16 illustrated as a portable computer that can include many or all of the elements and
17 features described herein relative to computer 602.

18 Logical connections between computer 602 and the remote computer 648
19 are depicted as a local area network (LAN) 650 and a general wide area network
20 (WAN) 652. Such networking environments are commonplace in offices,
21 enterprise-wide computer networks, intranets, and the Internet.

22 When implemented in a LAN networking environment, the computer 602 is
23 connected to a local network 650 via a network interface or adapter 654. When
24 implemented in a WAN networking environment, the computer 602 typically

1 includes a modem 656 or other means for establishing communications over the
2 wide network 652. The modem 656, which can be internal or external to computer
3 602, can be connected to the system bus 608 via the input/output interfaces 640 or
4 other appropriate mechanisms. It is to be appreciated that the illustrated network
5 connections are exemplary and that other means of establishing communication
6 link(s) between the computers 602 and 648 can be employed.

7 In a networked environment, such as that illustrated with computing
8 environment 600, program modules depicted relative to the computer 602, or
9 portions thereof, may be stored in a remote memory storage device. By way of
10 example, remote application programs 658 reside on a memory device of remote
11 computer 648. For purposes of illustration, application programs and other
12 executable program components such as the operating system are illustrated herein
13 as discrete blocks, although it is recognized that such programs and components
14 reside at various times in different storage components of the computing device
15 602, and are executed by the data processor(s) of the computer.

16 Computer 602 typically includes at least some form of computer readable
17 media. Computer readable media can be any available media that can be accessed
18 by computer 602. By way of example, and not limitation, computer readable
19 media may comprise computer storage media and communication media.
20 Computer storage media includes volatile and nonvolatile, removable and non-
21 removable media implemented in any method or technology for storage of
22 information such as computer readable instructions, data structures, program
23 modules or other data. Computer storage media includes, but is not limited to,
24 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
25

1 digital versatile discs (DVD) or other optical storage, magnetic cassettes, magnetic
2 tape, magnetic disk storage or other magnetic storage devices, or any other media
3 which can be used to store the desired information and which can be accessed by
4 computer 602. Communication media typically embodies computer readable
5 instructions, data structures, program modules or other data in a modulated data
6 signal such as a carrier wave or other transport mechanism and includes any
7 information delivery media. The term "modulated data signal" means a signal that
8 has one or more of its characteristics set or changed in such a manner as to encode
9 information in the signal. By way of example, and not limitation, communication
10 media includes wired media such as wired network or direct-wired connection,
11 and wireless media such as acoustic, RF, infrared and other wireless media.
12 Combinations of any of the above should also be included within the scope of
13 computer readable media.

14 The invention has been described herein in part in the general context of
15 computer-executable instructions, such as program modules, executed by one or
16 more computers or other devices. Generally, program modules include routines,
17 programs, objects, components, data structures, etc. that perform particular tasks
18 or implement particular abstract data types. Typically the functionality of the
19 program modules may be combined or distributed as desired in various
20 embodiments.

21 For purposes of illustration, programs and other executable program
22 components such as the operating system are illustrated herein as discrete blocks,
23 although it is recognized that such programs and components reside at various
24
25

1 times in different storage components of the computer, and are executed by the
2 data processor(s) of the computer.

3 Alternatively, the invention may be implemented in hardware or a
4 combination of hardware, software, and/or firmware. For example, one or more
5 application specific integrated circuits (ASICs) could be designed or programmed
6 to carry out the invention.

7

8 **Conclusion**

9 Thus, a security model a trusted environment has been described in which
10 secrets can be securely stored for trusted applications and in which the trusted
11 applications can be authenticated to third parties. These properties of the trusted
12 environment are maintained, even though various parts of the environment may be
13 upgraded or changed in a controlled way on the same computing device or
14 migrated to a different computing device.

15 Although the description above uses language that is specific to structural
16 features and/or methodological acts, it is to be understood that the invention
17 defined in the appended claims is not limited to the specific features or acts
18 described. Rather, the specific features and acts are disclosed as exemplary forms
19 of implementing the invention.